

10 GIGABIT ETHERNET MAPPINGS FOR A COMMON
LAN/WAN PMD INTERFACE WITH A SIMPLE UNIVERSAL
PHYSICAL MEDIUM DEPENDENT INTERFACE

CROSS-REFERENCE TO RELATED APPLICATIONS

- 5 This application is based on United States
Provisional Application No. 60/181,549, filed
February 10, 2000.

MICROFICHE APPENDIX

Not Applicable.

10 **TECHNICAL FIELD**

The present invention relates to high-speed data communications networks, and in particular to a method and apparatus for interfacing adjacent high-speed data communications networks.

15 **BACKGROUND OF THE INVENTION**

- Applicant's co-pending Canadian Patent Application Serial No. 2,273,522 proposes a technique of mapping high speed (HS) Ethernet frames (e.g. originating in a local area network [LAN] using a fiber distributed data interface [FDDI] physical layer) into a synchronous container (e.g. thin synchronous optical network [SONET] STS-192C or synchronous digital hierarchy [SDH] VC-4-64C frames) for transport across a synchronous packet switched network fabric at a standard OC-192 line rate of 9.95328 Gbaud, to
20 achieve a data rate of 9.5864Gb/s. This data rate is obtained independently of the LAN/MAN/WAN configuration of the network. The use of standard line rates in all configurations means that a common physical medium
25

00739385 121900

dependent (PMD) interface can be used, which reduces costs by permitting increased economies of scale. Additionally, conventional SONET/SDH routing systems (e.g. routers, cross-connects and add-drop multiplexors) can be used to transport data traffic in all configurations. Despite its versatility, the technique of Canadian Patent Application Serial No. 2,273,522 is unattractive for use in LAN configurations, because HS Ethernet LANs commonly provide a data rate of 10Gb/s. As a result, implementation of this technique in LAN configurations would have the effect of degrading overall data performance.

An alternative protocol has recently been proposed for mapping HS Ethernet frames into synchronous containers, that is intended to achieve a data rate of 10Gb/s for LAN configurations. However, in order to obtain this data rate, the synchronous container must be transmitted at an accelerated line rate of 10.3125 Gbaud. This non-standard line rate means that conventional SONET/SDH routing systems cannot readily be used in LAN configurations. Furthermore, the use of an accelerated line rate reduces data transmission reliability. According to this proposal, the standard OC-192 line rate of 9.95328 Gbaud is used for WAN configurations. However, the use of different line rates for LAN and WAN configurations means that respective different PMD interface devices must be used for each configuration. This latter requirement increases costs by reducing economies of scale.

Accordingly, a method and apparatus of mapping high-speed (HS) Ethernet signals to a synchronous container for transport across a synchronous packet-switched network fabric, which enables an Ethernet data rate of 10Gb/s in LAN configurations of the network, while using a standard

006672T 12141ROUS04U

OC-192 line rate of 9.95328 Gbaud, remains highly desirable.

SUMMARY OF THE INVENTION

5 An object of the present invention is to provide a method and apparatus for mapping Ethernet frames to a synchronous container such that an Ethernet data rate of 10Gb/s in LAN configurations can be obtained at a standard line rate of 9.95328 Gbaud.

10 Another object of the present invention is to provide a method and apparatus for mapping Ethernet frames to a synchronous container such that a common physical medium dependent (PMD) interface can be used independently of a LAN/WAN configuration of the network.

15 Accordingly, an aspect of the present invention provides a method of transporting a high speed Ethernet data stream comprising a sequential series of media access control (MAC) frames and having a data rate of 10Gb/s across a synchronous packet switched network fabric having a line rate of 9.953280 Gbaud. The method comprises the
20 steps of: at a sending interface: providing a synchronous container; compressing the Ethernet data stream to produce a compressed data stream; mapping the compressed data stream to the synchronous container; and launching the synchronous container across the network fabric. At a
25 receiving interface: the compressed data stream is extracted from the synchronous container; and the compressed data stream is decompressed to generate a recovered Ethernet data stream.

30 A further aspect of the present invention provides a method of interfacing a local area network (LAN) having

09739385 121900

006121-5866760

a 10Gb/s data rate and a synchronous packet switched physical network fabric having a line rate of 9.953280 Gbaud. The method comprises the steps of: receiving an Ethernet data stream comprising a sequential series of
5 media access control (MAC) frames; compressing the Ethernet data stream; and mapping the compressed Ethernet data stream to a synchronous container.

Another aspect of the present invention provides an interface for coupling a local area network (LAN) having
10 a 10Gb/s data rate to a synchronous packet switched physical network fabric having a line rate of 9.953280 Gbaud. The interface comprises: means for receiving an Ethernet data stream comprising a sequential series of media access control (MAC) frames; means for compressing
15 the Ethernet data stream; and means for mapping the compressed Ethernet data stream to a synchronous container.

In embodiments of the invention, the synchronous container comprises: a line/column frame format conforming to a conventional synchronous optical network (SONET) frame
20 format; and a stripped transport overhead (TOH). The frame format of the synchronous container may conform to a SONET synchronous transport signal-level 192 (STS-192) frame format. The stripped TOH may comprise only A1 and A2 octets of a conventional SONET frame TOH. In such cases,
25 the stripped TOH preferably comprises a predetermined number (e.g. at most 24) of each of the A1 and A2 octets.

In embodiments of the invention, the step of compressing the Ethernet data stream comprises removing an inter-frame gap (IFG) between successive MAC frames. Idle
30 MAC frames may also be removed from the Ethernet data stream.

The compressed data stream may be mapped to the synchronous container by adding successive frames of the compressed data stream directly to the SONET/SDH frame. In such cases, successive MAC frames of the compressed data stream are preferably added to the synchronous container starting immediately following the stripped TOH.

In some embodiments of the invention, the synchronous container may be launched into the network fabric by inverse multiplexing the synchronous container into a plurality of data streams, which are then launched into a respective plurality channels of the network fabric. In such cases, receiving the synchronous container includes receiving a respective data stream through each one of the plurality of channels. The received data streams are then aligned, and multiplexed to recover the synchronous container.

In embodiments of the invention, the compressed data stream is extracted from the synchronous container by: synchronizing a read clock with the synchronous container; and reading successive MAC frames of the compressed data stream from the synchronous container starting immediately following the stripped TOH. The read clock can be synchronized with the synchronous container by detecting a transition between the A1 and A2 octets.

25 The compressed data stream may be decompressed by
inserting an inter-frame gap (IFG) between successive MAC
frames. Insertion of an IFG between successive MAC frames
may be accomplished by: monitoring a preamble portion of a
first MAC frame; reading a length of a data portion of the
30 first MAC frame from the monitored preamble portion, to
identify a trailing byte of the first MAC frame; and

09/03/2015 11:00

inserting an idle MAC frame into the compressed stream immediately following the identified trailing byte.

An advantage of the present invention is that a common physical medium dependent (PMD) interface can be used, independently of the LAN/MAN/WAN configuration of the network, while ensuring a reliable 10Gb/s data rate in the LAN configuration.

BRIEF DESCRIPTION OF THE DRAWINGS

Further features and advantages of the present invention will become apparent from the following detailed description, taken in combination with the appended drawings, in which:

Fig. 1. is a block diagram illustrating a communications network including a pair of network nodes connected by a network medium, in which the present invention may be deployed;

Fig. 2 is a block diagram schematically illustrating a prior art physical layer interface;

Fig. 3 is a block diagram schematically illustrating a physical layer interface in accordance with an embodiment of the present invention;

Fig. 4 schematically illustrates encapsulation of Ethernet data streams in a synchronous container for transport through a WAN configured physical network fabric in accordance with an embodiment of the present invention;

Fig. 5 schematically illustrates encapsulation of Ethernet data streams in a synchronous container for transportation through a LAN configuration physical network

006121" SEE 6/60

fabric in accordance with an embodiment of the present invention;

Figs. 6a and 6b schematically illustrate compression and decompression of Ethernet data streams
5 usable in conjunction with the embodiment of Fig. 5;

Fig. 7 is a table showing a comparison between data and line rates obtained using the prior art interface of Fig. 2 and the present invention; and

Fig. 8 is a block diagram schematically
10 illustrating a physical layer interface in accordance with an alternative embodiment of the present invention.

It will be noted that throughout the appended drawings, like features are identified by like reference numerals.

15 **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT**

Fig. 1 illustrates a communications network including a pair of nodes 2a,2b (such as personal computers and LAN servers) connected for data communications through a network medium 4. The nodes 2a,2b are represented by the
20 standard Ethernet reference model which utilizes a multi-layer architecture made up of a physical layer (PHY) 6 and a data link layer 8. The PHY 6 is divided into three sub-layers, namely a physical coding sub-layer (PCS) 10, a physical medium attachment (PMA) sub-layer 12,
25 and a physical medium dependent (PMD) sub-layer 14. A Medium Dependent Interface (MDI) 16 couples the PHY 6 to the medium 4 (i.e. the network fabric). A Medium Independent Interface (MII) 18 couples the PHY 6 to the data link layer 8, which is composed of Media Access
30 Control (MAC) 20 and (optionally) MAC control 22

09739385 "121900

sub-layers. Upper layers 24 containing, for example, end-user software applications (not shown) are coupled to the data link layer 8 to facilitate communications between each node 2. This reference model renders the PHY 6 and medium 4 effectively transparent to upper-layer software applications, which are therefore enabled to communicate with each other using a common communications protocol that is independent of the medium 4. In Fig. 1, this is illustrated by a dashed line indicating a direct virtual connection 26 between applications in the upper layer 24. The actual flow of data between the upper layer applications follows a path 28 extending from the upper layer 24 of one node 2a, down through the data link layer 8 and PHY 6 of the device 2a, through the medium 4, and upwards through the PHY 6 and data link layer of the other node 2b.

In general, the medium 4 provides the network fabric, including physical links (not shown) and switching (and or routing) systems (not shown) which enable the physical transport of data between nodes 2a,2b connected to the network. The medium 4 may be configured as a wide area network (WAN) or a local area network (LAN), and may use any one of a variety of physical transport technologies (e.g. fiber-optics, copper, or twisted pair) and protocols (such as synchronous optical network [SONET], synchronous digital hierarchy [SDH], and Fiber Distributed Data Interface [FDDI]).

Local area networks (LANs) are commonly employed to connect general purpose computers (e.g. hosts or clients) and special purpose computers (e.g. servers) to provide communal file storage and access, e-mail services, etc. The oldest and best known LAN technology is the Ethernet,

which in many cases operates on a fiber distributed data interface (FDDI) PHY 6 using a common bus optical fiber medium 4.

5 A wide area network (WAN) is typically a packet switched data network that connects various packet switching nodes (e.g. routers, cross-connects and add-drop multiplexors) and transmission facilities to end-user communications devices (e.g. concentrator access switches or service provider routers). A WAN differs from a LAN in
10 geographic coverage, data rates, and also in data transport technology. In the current network space, a popular WAN technology is based on the SONET/SDH protocol using an optical fiber medium configured for Wave Division Multiplex (WDM) or Dense Wave Division Multiplex (DWDM) data
15 transport.

A metropolitan area network (MAN) is generally intended to serve network requirements that are geographically larger than a LAN, and smaller than a WAN. For example, a MAN may be used to provide a high band-width
20 backbone interconnecting distributed LANs. This type of application is typical of campus networks, where semi-autonomous LANs in multiple buildings are interconnected by a MAN to facilitate inter-building communications. In view of increasing complexity of
25 communications requirements, MANs are tending to evolve toward the adoption of WAN technology (e.g. SONET/SDH over DWDM optical fiber medium). Accordingly, for the purposes of the present application, MANs are considered to be a subset of WANs, and references to WANs herein should be
30 understood to also refer to MANs.

The desirability of using a packet switched routing technology (e.g. SONET/SDH over WDM or DWDM optical fiber

006434 5865760

medium) as the medium 4 for both LAN and WAN network configurations is well recognized. However, in a high speed data network environment, achievement of this goal is complicated by the fact that high speed Ethernet LANs typically achieve a 10Gb/s data rate (in the data link layer 8 and upper layer 24), whereas the packet switched data network operating under the SONET/SDH protocol uses an OC-192 line rate of 9.95328 Gbaud to obtain a payload data rate of 9.58464 Gb/s. In order to carry Ethernet traffic across a packet switched network medium, it is necessary to resolve this difference between the Ethernet data rate in the data link layer 8 and the payload rate within the medium 4.

Fig. 2 schematically shows a prior art physical layer interface for coupling high speed Ethernet traffic to a SONET/SDH packet switched network medium 4. As shown in Fig. 2, the conventional PCS 10 and PMA 12 are grouped together as a combined PCS/PMA HARI 30 which mediates data transport between the (conventional) data link layer 8 and a modified PMD sub-layer 14a. The HARI 30 exchanges data with the modified PMD 14a using 4 parallel data channels (in each direction) operating at a line rate of 3.125 Gbaud. The modified PMD 14a is subdivided into 8B/10B and 64/66 encoding layers 32,34; a framer FIFO 36; a scrambler 38; and a conventional electron/optical conversion layer 40. Thus in the prior art device of Fig. 2, much of the data encoding, framing, and scrambling functionality conventionally performed in the PCS 10 and PMA 12 are relocated into the PMD 14a. The PMD 14a is designed to operate in one of two modes depending on the LAN/WAN configuration of the network (see Fig. 7). In particular, in a LAN configuration, this interface is intended to achieve a data rate in the data

09739385 "121900

link layer 8 of 10Gb/s, in conformance with the high speed Ethernet standard. In order to obtain this data rate, the framer FIFO 36, scrambler 38 and converter 40 of the PMD 14a (and medium 4) is operated at a line rate
5 of 10.3125 Gbaud. This line rate is significantly higher than that supported by the SONET/SDH standard, which precludes the use of conventional (and legacy) SONET/SDH routing systems within the network medium 4.

In contrast, in a WAN configuration, the framer
10 FIFO 36, scrambler 38 and converter 40 of the PMD 14 and medium 4 are operated at the standard SONET/SDH line rate of 9.95328 Gbaud, which means that conventional and legacy SONET/SDH routing systems may be used in the medium 4. However, in this configuration, the data rate (in the data
15 link layer 8) is reduced to 9.29419Gb/s, which is somewhat lower than is normally expected based on the payload rate of conventional SONET/SDH signals. Additionally, the use of different line rates for LAN and WAN configuration implies that different devices must be used for the PMD
20 layer 14a in each configuration, which leads to increased costs.

The present invention provides a unified PHY 6b which allows the PMD 14 and medium 4 to operate at the standard OC-192 line rate of 9.95328 Gbaud while providing
25 a standard Ethernet data rate (in the data link layer 8) of 10Gb/s in LAN configurations and a SONET/SDH standard 9.58464Gb/s in WAN configurations. Features and operations of an exemplary embodiment of the present invention are described below with reference to Figs. 3
30 through 7. Alternative embodiments of the present invention will then be described with reference to Fig. 8.

Fig. 3 illustrates a unified PHY 6b in accordance with a first embodiment of the present invention. The unified PHY 6b is constructed using a layered model similar to the Ethernet reference model shown in Fig. 1. Thus the

5 PHY 6b is composed of a physical coding sub-layer 42, a physical medium attachment sub-layer 44, and a physical medium dependent sub-layer 46. The physical coding sub-layer 42 is coupled to the data link layer 8 via a conventional medium independent interface 18, which may

10 include one or more HARI interface devices 18a. Data transport through the medium independent interface 18 may be accomplished using a four lane data bus running at 3.125 Gbaud per lane and/or a thirty-two lane data bus running at 156.25MHz, to transfer an Ethernet data stream

15 consisting of MAC frames and delivering a data rate of 9.58464Gb/s in WAN configurations and 10Gb/s in LAN configurations. The operations of the data link layer 8 and MII 18 are known in the art and therefore will not be described in detail herein. Similarly, the PMD

20 sub-layer 46, which is coupled to the data transport medium 4 of the network, operates to perform electrical/optical signal conversion at an OC-192 line rate of 9.95328 Gbaud in a conventional manner, and thus will not be described in greater detail. The PCS 42 and the PMA

25 sub-layer 44 respectively operate to perform MAC packet delineation, and insertion/extraction of MAC packets into and from synchronous containers, to enable transport through the packet switched data network medium 4. These layers also function to perform a double scrambling of data

30 streams to improve system performance in a manner known in the art. Like the PMD sub-layer 46, both the PCS 42 and the PMA sub-layer 44 operate at a line rate of 9.95328 Gbaud.

006727 5855750

006721" 58E6E60

In accordance with the present invention, the manner in which Ethernet data streams are processed by the PCS 42, and the PMA sub-layer 44 is controlled based on the LAN or WAN configuration of the network. In particular, in a WAN configuration of the network, the PCS 42 performs signal scrambling in a conventional manner to produce a PCS data stream. The PMA sub-layer 44 then operates to map the PCS data stream into a conventional SONET synchronous payload envelope (SPE), which is inserted into a conventional synchronous container (e.g. an STS-192 frame) for transmission through the network medium 4 in a conventional manner. Conversely, in a LAN configuration of the network, the PCS 42 operates to produce the PCS data stream by compressing the Ethernet data stream, in addition to the conventional scrambling step. The PMA sub-layer 44 then maps this compressed PCS data stream into a stripped synchronous container for transport through the network medium 4. Selection of the mode of operation (e.g. for LAN or WAN network configurations) can be made by provisioning at the time of installation of the PHY 6b, or alternatively by means of an auto detection algorithm following power up of the PHY 6b. Operation of the PCS 42 and the PMA sub-layer 44 in each of these modes is described in greater detail below with reference to Figs. 4 through 7.

Fig. 4 illustrates the operation of the unified PHY 6a in a WAN configuration of the network. In this mode of operation, an outgoing Ethernet data stream is processed by the PCS 42 (e.g. $x^{43}+1$ scrambling) in a conventional manner to produce a PCS data stream 48 composed of a sequential series of MAC packets 50 separated by idle frames 52. The PCS data stream 48 is then forwarded to the PMA sub-layer 44 which performs a second scrambling of the PCS data stream 48 (e.g. with x^7+x^6+1) in a conventional

manner, before mapping the scrambled data stream into a synchronous payload envelope (SPE) 54. The SPE 54 is then inserted in a synchronous container 56 (e.g. an STS-192 frame) for transport through the network medium 4. As may be seen in Fig. 4, the construction of the synchronous container 56 and the SPE 54 conform to the SONET/SDH standard so that the synchronous container 56 can be transported through the network medium 4 using conventional (and legacy) SONET/SDH technology and transmission facilities.

The processing of synchronous containers 56 received by the PHY 6b is substantially the reverse of that described above. Thus the PMA sub-layer 44 processes each received synchronous container 56 to extract its payload data as a data stream. This data stream is then de-scrambled (e.g. using x^7+x^6+1) to recover the original PCS data stream 48, which is then passed to the PCS 42. The PCS 42 then processes the PCS data stream 48 by performing a second de-scrambling step (e.g. using $x^{43}+1$) and MAC packet delineation to generate a recovered Ethernet data stream which is passed to the data link layer 8 via the MII 18.

From the above description, it will be seen that for a WAN configuration of the network, the data throughput performance will be compatible with the payload rate of the packet switched data network. In the case of a SONET/SDH network operating at an OC-192 line rate of 9.95328 Gbaud, the synchronous container 56 can conveniently be an STS-192 frame, and the payload rate of the associated STS-192 SPE 54 will translate into a data rate of 9.58464Gb/s. As may be seen in Fig. 7, this data rate is higher than the

data rate of 9.29419Gb/s obtained by the prior art interface illustrated in Fig. 2 in WAN configurations.

5 Figs 5 and 6 illustrate operation of the PHY 6b in LAN configurations of the network. Fig. 5 illustrates a synchronous container 60 usable for transportation of payload data through the network medium 4 in this mode of operation. The synchronous container 60 is formed as a stripped STS frame, in that it retains the row and column format of a conventional STS frame. However, most of the
10 octets of the conventional transport overhead (TOH) have been removed. In the illustrated embodiment, only the A1 and A2 octets 62 of a conventional frame transport overhead are provided in the synchronous container 60. The A1 and A2 octets 62 are sufficient to enable routing of the
15 synchronous container 60 through the network medium 4, as well as synchronization and octet delineation of received containers. Because all routing through the network medium 4 is point-to-point in LAN configurations, all of the payload data inserted into a synchronous container 60
20 is destined for a single destination node on the network. Routing of individual MAC packets to MAC addresses subtending the destination node is handled by data link layer 8 processing within the destination node. Accordingly, the path overhead and fixed stuff normally
25 incorporated within an STS frame is not required. As a result, the use of a synchronous payload envelope is also not required for this mode of operation.

30 In order to maximize utilization of the synchronous container 60, the compressed data stream received from the physical coding sub-layer 42 is mapped directly into the synchronous container 60, starting at the first octet following the A1 and A2 octets 62, without first being

006721"SEE/60

inserted into a Synchronous Payload Envelope (SPE). As a result, space normally used for TOH and Path overhead (POH) within a conventional STS-192c frame is freed-up for carrying the compressed data stream.

5 A further increase in the data rate is obtained by compressing the Ethernet data stream within the PCS 42 (in addition to the conventional scrambling step) to generate a compressed PCS data stream. In general, this can be accomplished by processing the Ethernet data stream to
10 eliminate the inter-frame gap 64 which normally separates successive MAC packets 66. As a result, MAC packets are inserted (following scrambling) into the synchronous container 60 back-to-back, with no dead space between the last bit of one packet 66a and a first bit of a successive
15 packet 66b. A further compression of the PCS data stream can be accomplished by processing the data stream to eliminate idle MAC packets (i.e. MAC packets in which the data field is empty) so that each MAC packet 66 inserted into the synchronous container 60 includes payload data.
20 Following insertion of the compressed PCS data stream into the synchronous container 60, the synchronous container 60 can be transported through the network medium 4 at a conventional OC-192 line rate of 9.95328 Gbaud. However, by maximizing utilization of the synchronous container 60,
25 coupled with compression of the Ethernet data stream, a data rate of 10Gb/s (in the data link layer 8) can be achieved. As can be seen in Fig. 7, this data performance compares favorably with the prior art interface shown in Fig. 2, which required a line rate of 10.3125 Gbaud to
30 obtain a data rate of 10Gb/s in LAN configurations.

Processing of received synchronous containers to obtain a recovered Ethernet data stream is substantially

006121 "3866250

the reverse of that described above. Thus, upon reception of a synchronous container 60, the PMA layer 44 extracts the payload data from the synchronous container 60 as a serial data stream starting immediately following the A1 and A2 octets 62 of the container 60. The serial data stream is then de-scrambled (e.g. using x^7+x^6+1) and passed to the PCS 42 as a recovered compressed PCS data stream. Within the PCS 42 the compressed PCS data stream is decompressed by insertion of an inter-frame gap (of, for example, 11 octets) between successive MAC frames to produce a regenerated PCS data stream which can be de-scrambled a second time (e.g. using $x^{43}+1$) to produce the regenerated Ethernet data stream.

As shown in Fig. 6, compression and decompression of the PCS data stream can be accomplished by examining the length field 68 of the preamble of each MAC packet 66, which is conventionally used to specify the length of the data field 70 of the respective MAC packet 66. Thus idle packets and inter-frame gaps can be identified by a zero value in the length field 68, thereby allowing such idle packets and inter-frame gaps to be dropped from the PCS data stream. Once this has been accomplished, the compressed PCS data stream comprises a serial sequence of data-carrying MAC packets 66 lying back-to-back. This stream can be subsequently decompressed by again examining the length field 68 of the MAC packet preamble. Because the preamble is of a known predetermined length, the value contained in the length field 68 can be used as a pointer to the first bit of the next successive MAC packet 66. Thus as the regenerated compressed PCS data stream is received by the PCS 42, MAC packet delineation can be performed by examining the length field 68 of each successive MAC packet 66 to locate the first bit of the

006727 58E6E760

next packet 66. The data stream can then be further processed to insert an inter-frame gap of a predetermined length between successive packets 66. Preferably, the inserted inter-frame gap has a length of not more than eleven octets.

Fig. 8 shows an alternative embodiment of the interface of the present invention. In this case, the physical coding sub-layer 42 and physical medium attachment sub-layer 44 are implemented within a single device. The physical medium dependent sub-layer 46 is implemented as a second device, and a communications bus 70 is provided to interconnect these two devices. In one embodiment, the communications bus 70 comprises a multi-lane parallel bus running at a standard lower-speed SONET lint rate. In the illustrated embodiment, the communications bus 70 is implemented as a four-lane parallel bus running at a standard SONET STS-48 line rate of 2.48832 Gbaud. This bus 70 can be implemented using copper connections with a reach of up to twenty inches, without significantly degrading data transmission reliability. In use, the synchronous container 60 can be distributed (i.e. inverse multiplexed using, for example a conventional bit-slicing technique) across each of the lanes of the communications bus 70 by the physical medium attachment sub-layer 44 and transmitted to the physical medium dependent layer 46 for transport across the network medium 4. If desired, the multi-lane data transmission bus 70 can be propagated through the network medium 4 by transmitting the data streams in each of the lanes through the network medium 4 using respective different channels (e.g. wavelengths or fibers) running at a suitable conventional line rate (e.g. OC-48 in the case of the four-lane communications bus 70 illustrated in FIG. 8). In this case, the A1 and A2 octets

of the synchronous container can be used to realign the respective data streams received through each of the channels at the destination node. In particular, by bit-slicing the synchronous container across multiple
5 lanes, the respective data stream transmitted/received through each channel will each contain an A1/A2 octet transition, which will be at an identical position within each data stream. Thus a receiving physical medium dependent layer 46 can use the A1/A2 transition to bit-
10 align the received data streams with each other using techniques known in the art, such as, for example, by buffering each data stream in a respective elastic store (not shown) and adjusting a read pointer based on the detected location of the A1/A2 transition within the
15 respective data stream. The thus bit-wise aligned data streams can then be multiplexed (e.g. using a conventional bit-interleaving technique) to recover the original synchronous container.

In an alternative embodiment, a conventional 16 bit
20 low voltage differential signal (LVDS) chip interface can be used between the physical medium dependent layer 46 and the physical medium attachment layer 44.

Based on the foregoing, it will be seen that the present invention provides a technique for transmitting
25 high speed Ethernet data streams across a packet switched physical network medium 4. The present invention utilizes a common physical medium dependent interface 46 that can be used independently of the LAN/WAN configuration of the network. Additionally, data rates compatible with
30 conventional SONET payload rates are obtained in a WAN configuration, while a data rate of 10Gb/s is obtained in the LAN configuration.

006121-5866260

